arXiv:1411.2536v2 [cs.DC] 1 Apr 2015

Algorithmic Energy Saving for Parallel Cholesky, LU, and QR Factorizations

Li Tan*, Zizhong Chen

Department of Computer Science and Engineering, University of California, Riverside, CA, USA

Abstract

The pressing demands of improving energy efficiency for high performance scientific computing have motivated a large body of software-controlled hardware solutions using Dynamic Voltage and Frequency Scaling (DVFS) that strategically switch processors to low-power states, when the peak processor performance is not necessary. Although OS level solutions have demonstrated the effectiveness of saving energy in a black-box fashion, for applications with variable execution characteristics, the optimal energy efficiency can be blundered away due to defective prediction mechanism and untapped load imbalance. In this paper, we propose `TX`, a library level *race-to-halt* DVFS scheduling approach that analyzes Task Dependency Set of each task in parallel Cholesky, LU, and QR factorizations to achieve substantial energy savings OS level solutions cannot fulfill. Partially giving up the generality of OS level solutions per requiring library level source modification, `TX` leverages algorithmic characteristics of the applications to gain greater energy savings. Experimental results on two power-aware clusters indicate that `TX` can save up to 17.8% more energy than state-of-the-art OS level solutions with negligible 3.5% on average performance loss.

Keywords: energy, DVFS, library level, algorithmic characteristics, task dependency set, critical path, parallel Cholesky, LU, and QR factorizations

*Corresponding Author: darkwhite29@gmail.com

1. Introduction

1.1. Motivation

With the growing prevalence of distributed-memory architectures, high performance scientific computing has been widely employed on supercomputers around the world ranked by the TOP500 list [1]. Considering the crucial fact that the costs of powering a supercomputer are rapidly increasing nowadays due to expansion of its size and duration in use, improving energy efficiency of high performance scientific applications has been regarded as a pressing issue to solve. The Green500 list [2], ranks the top 500 supercomputers worldwide by performance-power ratio in six-month cycles. Consequently, root causes of high energy consumption while achieving performance efficiency in parallelism have been widely studied. With different focuses of studying, holistic hardware and software approaches for reducing energy costs of running high performance scientific applications have been extensively proposed. Software-controlled hardware solutions such as DVFS-directed (Dynamic Voltage and Frequency Scaling) energy efficient scheduling are deemed to be effective and lightweight [3] [4] [5] [6] [7] [8] [9]. Performance and memory constraints have been considered as trade-offs for energy savings [10] [11] [12] [13] [14].

DVFS is a runtime technique that is able to switch operating frequency and supply voltage of a hardware component (CPU, GPU, memory, etc.) to different *levels* (also known as *gears* or *operating points*) per workload characteristics of applications to gain energy savings dynamically. CPU and GPU are the most widely applied hardware components for saving energy via DVFS due to two primary reasons: (a) Compared to other components such as memory, CPU/GPU DVFS is easier to implement [15] – various handy DVFS APIs have been industrialized for CPU/GPU DVFS such as CPUFreq kernel infrastructure [16] incorporated into the Linux kernel and NVIDIA System Management Interface (nvidia-smi) [17] for NVIDIA GPU; (b) CPU energy costs dominate the total system energy consumption [18] (CPU and GPU energy costs dominate if heterogeneous architectures are considered), and thus saving CPU and GPU energy greatly improves energy efficiency of the whole system. In this work, we focus on distributed-memory systems without GPU. For instance, energy saving opportunities can be exploited by reducing CPU frequency and voltage during non-CPU-bound operations such as large-message MPI communication, since generally execution time of such operations barely increases at a low-power state of CPU, in con-

trast to original runs at a high-power state. Given the fact that energy consumption equals product of average power consumption and execution time ($E = \overline{P} \times T$), and the assumption that dynamic power consumption of a CMOS-based processor is proportional to product of operating frequency and square of supply voltage ($P \propto fV^2$) [19] [20], energy savings can be effectively achieved using DVFS-directed strategical scheduling approaches with little performance loss.

Running on distributed-memory architectures, HPC applications can be organized and scheduled in the unit of task, a set of operations that are functionally executed as a whole. Different tasks within one process or across multiple processes may depend on each other due to *intra-process* and *inter-process* data dependencies. Parallelism of task-parallel algorithms and applications can be characterized by graph representations such as Directed Acyclic Graph (DAG), where data dependencies among parallel tasks are appropriately denoted by directed edges. DAG can be effective for analyzing parallelism present in HPC runs, which is greatly beneficial to achieving energy efficiency. As typical task-parallel algorithms for scientific computing, dense matrix factorizations in numerical linear algebra such as Cholesky, LU, and QR factorizations have been widely adopted to solve systems of linear equations. Empirically, as standard functionality, routines of dense matrix factorizations are provided by various software libraries of numerical linear algebra for distributed-memory multicore architectures such as ScaLAPACK [21], DPLASMA [22], and MAGMA [23]. Therefore, saving energy for parallel Cholesky, LU, and QR factorizations contributes significantly to the greenness of high performance scientific computing nowadays.

1.2. Limitations of Existing Solutions

Most existing energy saving solutions for high performance scientific applications are (combination of) variants of two classic approaches: (a) A Scheduled Communication (SC) approach [3] [24] [25] [8] that keeps low CPU performance during communication and high CPU performance during computation, as large-message communication is not CPU-bound while computation is, and (b) a Critical Path (CP) approach [8] [9] [26] [27] that guarantees that tasks on the CP run at the highest CPU frequency while reduces frequency *appropriately* (i.e., without further delay to incur overall performance loss) for tasks off the CP to minimize slack.

Per the operating layer, existing solutions can be categorized into two types: OS level and application level. In general, OS level solutions feature

two properties: (a) Working aside running applications and thus requiring no application-specific knowledge and no source modification, and (b) making online energy efficient scheduling decisions via dynamic monitoring and analysis. However, application level solutions statically utilize application-specific knowledge to perform specialized scheduling for saving energy, generally with source modification and recompilation (i.e., *generality*) trade-offs. Although with high generality, OS level solutions may suffer from critical disadvantages as follows, and consequently are far from a sound and complete solution, for applications such as parallel Cholesky, LU, and QR factorizations in particular:

Effectiveness. Although intended to be effective for general applications, OS level approaches rely heavily on the underlying workload prediction mechanism, due to lack of knowledge of application characteristics. One prediction mechanism can work well for a specific type of applications sharing similar characteristics, but can be error-prone for other applications, in particular, applications with random/variable execution characteristics where the prediction mechanism performs poorly. Algorithms presented in [25] [8] [9] predict execution characteristics of the upcoming interval (i.e., a fixed time slice) according to recent intervals. This prediction mechanism is based on a simple assumption that task behavior is identical every time a task is executed [9]. However, it can be defective for applications with random/variable execution characteristics, such as matrix factorizations, where the remaining unfinished matrices become smaller as the factorizations proceed. In other words, length variation of iterations of the core loop for matrix factorizations can make the above prediction mechanism inaccurate, which invalidates potential energy savings. Moreover, since new technologies such as Hyper-Threading [28] have emerged for exploiting thread level parallelism on distributed-memory multicore systems, e.g., MPI programs can be parallelized on local multicore processors using OpenMP, behavior among parallel tasks can vary due to more non-deterministic events occurred in the multithreaded environment.

Further, the OS level prediction can be costly and thus energy savings are diminished. Recall that OS level solutions must predict execution details in the next interval using prior execution information. However, execution history in some cases may not necessarily be a reliable source for workload prediction, e.g., for applications with fluctuating runtime patterns at the beginning of the execution. As such, it can be time-consuming for obtaining accurate prediction results. Since during the prediction phase no energy

savings can be fulfilled, considerable potential energy savings can be wasted for accurate but lengthy prediction.

Completeness. OS level solutions only work when tasks such as computation and communication are being executed, energy saving opportunities are untapped during the time otherwise. Empirically, even though load balancing techniques have been leveraged, due to data dependencies among tasks and load imbalance that is not completely eliminated, not all tasks in different processes across nodes can start to work and finish at the same time. More energy can be saved for tasks waiting at the beginning of an execution, and for the last task of one process finishing earlier than that of other processes across nodes (details are illustrated in Figure 4). Restricted by the daemon-based nature of working aside real running tasks, OS level solutions cannot attain energy savings for such tasks.

1.3. Our Contributions

In this paper, we propose a library level *race-to-halt* DVFS scheduling approach via Task Dependency Set (TDS) analysis based on algorithmic characteristics, namely `TX` (TDS-based race-to-halt), to save energy for task-parallel applications running on distributed-memory architectures, taking parallel Cholesky, LU, and QR factorizations for example. The idea of library level *race-to-halt* scheduling is intended for any task-parallel programming models where data flow analysis can be applied. The use of TDS analysis as a compiler technique allows possible extension of this work to a general compiler-based approach based on static analysis only. In summary, the contributions of this paper are as follows:

- Compared to application level solutions, for widely used software libraries such as numerical linear algebra libraries, `TX` restricts source modification and recompilation at library level, and replacement of the energy efficient version of the libraries is allowed at link time (i.e., with *partial loss of generality*). No further source modification and recompilation are needed for applications where the libraries are called;

- Compared to OS level solutions, `TX` is able to achieve substantial energy savings for parallel Cholesky, LU, and QR factorizations (i.e., with *higher energy efficiency*), since via algorithmic TDS analysis, `TX` circumvents the defective prediction mechanism at OS level, and also manages to save more energy from possible load imbalance;

- We formally model that `TX` is comparable to the *CP* approach in energy saving capability under the circumstance of current CMOS technologies that allow insignificant variation of supply voltage as operating frequency scales via DVFS;

- With negligible performance loss (3.5% on average), on two power-aware clusters, `TX` is evaluated to achieve up to 33.8% energy savings compared to original runs of different-scale matrix factorizations, and save up to 17.8% and 15.9% more energy than state-of-the-art OS level *SC* and *CP* approaches, respectively.

The rest of this paper is organized as follows. Section 2 introduces basics of parallel Cholesky, LU, and QR factorizations. We present TDS and CP in section 3, and our `TX` approach in section 4. Implementation details and experimental results are provided in section 5. Section 6 discusses related work and section 7 concludes.

2. Background: Parallel Cholesky, LU, and QR Factorizations

As classic dense numerical linear algebra operations for solving systems of linear equations, such as $Ax = b$ where A is a given coefficient matrix and b is a given vector, Cholesky factorization applies to the case that A is a symmetric positive definite matrix, while LU and QR factorizations apply to any general $M \times N$ matrices. The goal of these operations is to factorize A into the form LL^T where L is lower triangular and L^T is the transpose of L, the form LU where L is unit lower triangular and U is upper triangular, and the form QR where Q is orthogonal and R is upper triangular, respectively. Thus from $LL^Tx = b$, $LUx = b$, $QRx = b$, x can be easily solved via forward substitution and back substitution. In practice, parallel Cholesky, LU, and QR factorizations are widely employed in extensive areas of high performance scientific computing. Various software libraries of numerical linear algebra for distributed multicore scientific computing such as ScaLAPACK [21], DPLASMA [22], and MAGMA [23] provide routines of these matrix factorizations as standard functionality. In this section, we present basics of parallel Cholesky, LU, and QR factorizations, and introduce an effective graph representation for parallelism present during the execution of these matrix factorizations.

2.1. 2-D Block Cyclic Data Distribution

Regardless of the fact that nowadays matrix involved in many parallel numerical linear algebra operations is too large to fit into the memory of single node, the essence of a parallel algorithm is how it partitions workloads into a cluster of computing nodes as balanced as possible to better exploit parallelism, which is also referred to as load balancing. In numerical linear algebra operations, distributing a global matrix into a process grid in a linear fashion does not benefit a lot from parallelism, since although the data allocated into each computing node are balanced in terms of amount, parallel execution of computation and communication are restricted by frequently arising data dependencies among tasks performed by different processes on different nodes.

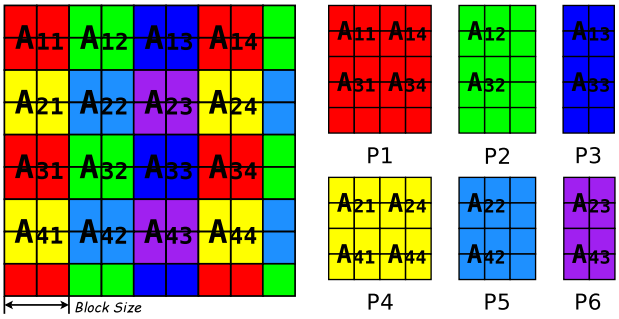

(a) Global View (b) Local View

Figure 1: 2-D Block Cyclic Data Distribution on a 2×3 Process Grid in Global View and Local View.

As shown in Figure 1, *2-D block cyclic data distribution*, an effective way for load balancing, has been widely used in various numerical linear algebra libraries, such as HPL [29], ScaLAPACK [21], DPLASMA [22], and MAGMA [23]. Specifically, the global matrix is partitioned in a two-dimensional block cyclic fashion, and blocks are mapped into different nodes cyclically along both rows and columns of the process grid, so that tasks without data dependencies are able to be executed by multiple nodes simultaneously to achieve parallelism. Figure 1 (a) shows the partitioned global matrix in a global view, while Figure 1 (b) depicts local matrices residing in different nodes individually, where the global matrix elements on one node are accessed periodically throughout the execution to balance the workloads in different nodes, instead of all at once as in the linear data distribution. The granularity of

partitioning is determined by a *block size*, either specified by the user, or automatically tuned according to hardware configuration or set by default by the application itself. In practice, the global matrix size may not necessarily be multiples of the chosen block size. Typically, due to uneven division, the existence of remainder matrices, i.e., the last row and the last column in Figure 1 (a), barely affects parallelism empirically given a fine-grained partition.

2.2. DAG Representation of Parallel Cholesky, LU, and QR Factorizations

A well-designed partitioning and highly-efficient parallel algorithms of computation and communication substantially determine energy and performance efficiency of task-parallel applications. For such purposes, classic implementations of parallel Cholesky, LU, and QR factorizations are as follows: (a) Partition the global matrix into a cluster of computing nodes as a process grid using 2-D block cyclic data distribution [30] for load balancing; (b) perform local *diagonal matrix* factorizations in each node individually and communicate factorized local matrices to other nodes for *panel matrix* solving and *trailing matrix* updating, as shown in Figure 2, a stepwise LU factorization without pivoting. Due to frequently-arising data dependencies, parallel execution of parallel Cholesky, LU, and QR factorizations can be characterized using Directed Acyclic Graph (DAG), where data dependencies among parallel tasks are appropriately represented. DAG for parallel Cholesky, LU, and QR factorizations is formally defined below:

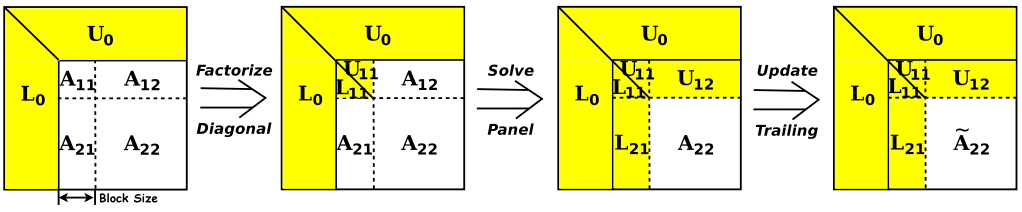

Figure 2: Stepwise Illustration of LU Factorization without Pivoting.

Definition 1. Data dependencies among parallel tasks of parallel Cholesky, LU, and QR factorizations running on a distributed-memory computing system are modeled by a Directed Acyclic Graph (DAG) $G = (V, E)$, where each node $v \in V$ denotes a task of Cholesky/LU/QR factorization, and each directed edge $e \in E$ represents a dynamic data dependency from task t_j to

$$\begin{pmatrix} A_{11} & A_{21}^T & A_{31}^T & A_{41}^T \\ A_{21} & A_{22} & A_{32}^T & A_{42}^T \\ A_{31} & A_{32} & A_{33} & A_{43}^T \\ A_{41} & A_{42} & A_{43} & A_{44} \end{pmatrix} = \begin{pmatrix} L_{11} & 0 & 0 & 0 \\ L_{21} & L_{22} & 0 & 0 \\ L_{31} & L_{32} & L_{33} & 0 \\ L_{41} & L_{42} & L_{43} & L_{44} \end{pmatrix} \times \begin{pmatrix} L_{11}^T & L_{21}^T & L_{31}^T & L_{41}^T \\ 0 & L_{22}^T & L_{32}^T & L_{42}^T \\ 0 & 0 & L_{33}^T & L_{43}^T \\ 0 & 0 & 0 & L_{44}^T \end{pmatrix}$$

$$= \begin{pmatrix} L_{11}L_{11}^T & L_{11}L_{21}^T & L_{11}L_{31}^T & L_{11}L_{41}^T \\ L_{21}L_{11}^T & L_{21}L_{21}^T + L_{22}L_{22}^T & L_{21}L_{31}^T + L_{22}L_{32}^T & L_{21}L_{41}^T + L_{22}L_{42}^T \\ L_{31}L_{11}^T & L_{31}L_{21}^T + L_{32}L_{22}^T & L_{31}L_{31}^T + L_{32}L_{32}^T + L_{33}L_{33}^T & L_{31}L_{41}^T + L_{32}L_{42}^T + L_{33}L_{43}^T \\ L_{41}L_{11}^T & L_{41}L_{21}^T + L_{42}L_{22}^T & L_{41}L_{31}^T + L_{42}L_{32}^T + L_{43}L_{33}^T & L_{41}L_{41}^T + L_{42}L_{42}^T + L_{43}L_{43}^T + L_{44}L_{44}^T \end{pmatrix}$$

Figure 3: Matrix Representation of a 4×4 Blocked Cholesky Factorization (We henceforth take parallel Cholesky factorization for example due to algorithmic similarity among three types of matrix factorizations).

task t_i that both tasks manipulate on either different *intra-process* or *inter-process* local matrices (i.e., an *explicit* dependency) or the same *intra-process* local matrix (i.e., an *implicit* dependency), denoted by $t_i \rightarrow t_j$.

Example. Due to similarity among the three matrix factorizations and space limitation, we henceforth take parallel Cholesky factorization for example to elaborate our approach. Consider a 4×4 blocked Cholesky factorization as given in Figure 3. The outcome of the task factorizing A_{11}, i.e., L_{11}, is used in the tasks solving local matrices L_{21}, L_{31}, and L_{41} in the same column as L_{11}, i.e., the tasks calculating the *panel matrix*. In other words, there exist three data dependencies from the tasks solving L_{21}, L_{31}, and L_{41} to the task factorizing A_{11}, denoted by three *solid* directed edges from the task Factorize(1,1) to the tasks Solve(2,1), Solve(3,1), and Solve(4,1) individually as shown in Figure 4 (see page 15). Besides the above *explicit* dependencies, there exists an *implicit* dependency between the task updating local matrix A_{32} and the task subsequently solving L_{32} on the same local matrix, denoted by the *dashed* directed edge from the task Update1(3,2) to the task Solve(3,2) in Figure 4. Note that communication among tasks is not shown in Figure 4, and updating diagonal local matrices and updating non-diagonal local matrices are distinguished as Update2() and Update1() respectively due to different computation time complexity.

3. Fundamentals: Task Dependency Set and Critical Path

Based on the task-parallel DAG representation, next we present Task Dependency Set (TDS) and Critical Path (CP) of running parallel Cholesky, LU, and QR factorizations, where TDS contains static dependency informa-

Table 1: Notation in Algorithms 1, 2, 3, and 4 and Henceforth.

$task, t_1, t_2$	One task of matrix factorizations
N_{proc}	Square root of the total number of processes
f_l	The lowest CPU frequency set by DVFS
f_h	The highest CPU frequency set by DVFS
f_{opt}	Optimal ideal frequency to finish a task with its slack eliminated
$ratio$	Ratio between split durations for optimal frequency approximation
$\mathsf{TDS}_{in}(task)$	TDS consisting of tasks that are depended by $task$ as the input
$\mathsf{TDS}_{out}(task)$	TDS consisting of tasks that depend on $task$ as the input
$CritPath$	One task trace to finish matrix factorizations with zero total slack
$slack$	Time that a task can be delayed by with no overall performance loss
$CurFreq$	Current CPU frequency in use
$DoneFlag$	Indicator of the finish of a task

tion of parallel tasks to utilize at runtime, and CP pinpoints potential energy saving opportunities in terms of slack among the tasks.

3.1. Task Dependency Set

For determining the appropriate timing for exploiting potential energy saving opportunities via DVFS, we leverage TDS analysis in our `TX` approach. Next we first formally define TDS, and then showcase how to generate two types of TDS for each task in Cholesky factorization using Algorithm 1. Producing TDS for LU and QR factorizations is similar with minor changes in Algorithm 1 per algorithmic characteristics. Table 1 lists the notation used in the algorithms and discussion in sections 3 and 4.

Definition 2. Given a task t of a parallel Cholesky/LU/QR factorization, data dependencies related to a data block manipulated by the task t are classified as elements of two types of TDS: $\mathsf{TDS}_{in}(t)$ and $\mathsf{TDS}_{out}(t)$, where dependencies from the data block to other tasks t_i are categorized into $\mathsf{TDS}_{out}(t)$ and denoted as t_i for short, and dependencies from other tasks t_j to the data block are categorized into $\mathsf{TDS}_{in}(t)$ and denoted as t_j for short.

Example. Consider the Cholesky factorization in Figure 3. Two TDS of each task can be generated statically per algorithmic characteristics as shown in Algorithm 1: Since the resulting local matrices of factorization tasks (e.g., L_{11}) are used in column-wise *panel matrix* solving (e.g., solving L_{21}, L_{31}, and L_{41}), data dependencies from *panel matrices* to factorized *diagonal matrices* are included in TDS_{in} of tasks solving *panel matrices* (e.g., $\mathsf{TDS}_{in}(\mathsf{S}(2,1))$,

Algorithm 1 *Task Dependency Set Generation Algorithm*

GenTDS($task$, N_{proc})

```
 1: switch (task)
 2:   case Factorize:
 3:     foreach i < j, 1 ≤ i, j ≤ N_proc do
 4:       insert(TDS_in(S(j, i)), F(i, i));
 5:       insert(TDS_out(F(i, i)), S(j, i));
 6:   case Update1:
 7:     foreach i < j, 1 ≤ i, j ≤ N_proc do
 8:       if (IsLastInstance(U_1(j, i))) then
 9:         insert(TDS_in(U_1(j, i)), S(j, i));
10:         insert(TDS_out(S(j, i)), U_1(j, i));
11:   case Update2:
12:     foreach 1 ≤ i ≤ N_proc do
13:       if (IsLastInstance(U_2(i, i))) then
14:         insert(TDS_in(F(i, i)), U_2(i, i));
15:         insert(TDS_out(U_2(i, i)), F(i, i));
16:   case Solve:
17:     foreach 1 ≤ i < j ≤ N_proc do
18:       foreach j < k ≤ N_proc do
19:         insert(TDS_in(U_1(k, j)), S(j, i));
20:         insert(TDS_out(S(j, i)), U_1(k, j));
21:       foreach i < k < j ≤ N_proc do
22:         insert(TDS_in(U_1(j, k)), S(j, i));
23:         insert(TDS_out(S(j, i)), U_1(j, k));
24:       insert(TDS_in(U_2(j, j)), S(j, i));
25:       insert(TDS_out(S(j, i)), U_2(j, j));
26: end switch
```

$\mathsf{TDS}_{in}(\mathsf{S}(3,1))$, and $\mathsf{TDS}_{in}(\mathsf{S}(4,1))$), and TDS_{out} of tasks factorizing *diagonal matrices* (e.g., $\mathsf{TDS}_{out}(\mathsf{F}(1,1))$). Likewise TDS_{in} and TDS_{out} of other tasks holding different dependencies can be produced following Algorithm 1.

3.2. Critical Path

Although load balancing techniques are leveraged for distributing workloads into a cluster of computing nodes as evenly as possible, assuming that all nodes have the same hardware configuration and thus the same computation and communication capability, slack can result from the fact that different processes can be utilized unfairly due to three primary reasons: (a) imbalanced computation delay due to data dependencies among tasks, (b) imbalanced task partitioning, and (c) imbalanced communication delay. Difference in CPU utilization results in different amount of computation slack. For instance, constrained by data dependencies, the start time of processes running on different nodes differs from each other, as shown in Figure 4 (see page 15) where $\mathsf{P1}$ starts earlier than the other three processes. Moreover, since the location of local matrices in the global matrix determines what types of computation are performed on the local matrices, load imbalancing from difference in task types and task amount allocated to different processes cannot be eliminated completely by the 2-D block cyclic data distribution, as shown in Figure 4 where $\mathsf{P2}$ has lighter workloads compared to the other three processes. Imbalanced communication time due to different task amount among the processes further extends the difference in slack length for different processes.

Critical Path (CP) is one particular task trace from the beginning task of one run of a task-parallel application to the ending one with the total slack of zero. Any delay on tasks on the CP increases the total execution time of the application, while dilating tasks off the CP into their slack individually without further delay does not cause performance loss as a whole. Energy savings can be achieved by appropriately reducing frequency to dilate tasks off the CP into their slack as much as possible, which is referred to as the *CP* approach. Numerous existing OS level solutions effectively save energy via *CP-aware* analysis [3] [25] [8] [9] [26] [27]. Figure 4 highlights one CP for the provided parallel Cholesky factorization with bold edges. We next present a feasible algorithm to generate a CP in parallel Cholesky, LU, and QR factorizations via TDS analysis.

Algorithm 2 *Critical Path Generation Algorithm via TDS Analysis*

```
GenCritPath(CritPath, task, N_proc)
 1: CritPath ← ∅
 2: switch (task)
 3:    case Factorize:
 4:       insert(CritPath, F(i, i))
 5:    case Update1:
 6:       Do Nothing
 7:    case Update2:
 8:       if (t_1 ∈ CritPath && t_1 ∈ TDS_out(U_2(i, i))) then
 9:          insert(CritPath, U_2(i, i))
10:    case Solve:
11:       foreach 1 ≤ i < j ≤ N_proc do
12:          if (t_2 ∈ CritPath && t_2 ∈ TDS_out(S(j, i))) then
13:             insert(CritPath, S(j, i))
14: end switch
```

3.3. Critical Path Generation via TDS

We can generate a CP for parallel Cholesky, LU, and QR factorizations as the basis of the *CP* approach using Algorithm 2. Consider the same parallel Cholesky factorization above. The heuristic of the CP generation algorithm is as follows: (a) Each task of factorizing is included in the CP, since the local matrices to factorize are always updated last, compared to other local matrices in the same row of the global matrix, and the outcome of factorizing is required in future computation. In other words, the task of factorizing serves as a transitive step that cannot be executed in parallel together with other tasks; (b) each task of Update1() is excluded from the CP, since it does not have direct dependency relationship with any tasks of factorizing, which are already included in the CP; (c) regarding Update2(), we select the ones that are directly depended by the tasks of factorizing on the same local matrix into the CP; (d) we choose the tasks of solving that are directly depended by Update2() (or directly depends on Factorize(), not shown in the algorithm) into the CP. Note that CP can also be identified using methods other than TDS analysis [3] [26] [27].

4. TX: Energy Efficient Race-to-halt DVFS Scheduling

In this section, we present in detail three energy efficient DVFS scheduling approaches for parallel Cholesky, LU, and QR factorizations individually: the *SC* approach, the *CP* approach, and our `TX` approach. We further demonstrate that `TX` is able to save energy substantially compared to the other two solutions, since via *TDS-based race-to-halt*, it circumvents the defective prediction mechanism employed by the *CP* approach at OS level, and further saves energy from possible load imbalance. Moreover, we formally prove that `TX` is comparable to the *CP* approach in energy saving capability under the circumstance of current CMOS technologies.

4.1. Custom Functions

In Algorithms 1, 2, 3, and 4, nine custom functions are introduced for readability: insert(TDS(t_1), t_2), delete(TDS(t_1), t_2), SetFreq(), IsLastInstance(), Send(), Recv(), IsFinished(), GetSlack(), and GetOptFreq(). The implementation of insert() and delete() is straightforward: Add task t_2 into the TDS of task t_1, and remove t_2 from the TDS of t_1. SetFreq() is a wrapper of handy DVFS APIs that set specific frequencies, and Send() and Recv() are wrappers of MPI communication routines that send and receive flag messages among tasks respectively. IsLastInstance() is employed to determine if the current task is the last instance of the same type of tasks manipulating the same data block, and IsFinished() is employed to determine if the current task is finished: Both are easy to implement at library level. GetSlack() and GetOptFreq() are used to get slack of a task, and calculate the optimal CPU frequency to dilate a task into its slack as much as possible, respectively. Implementing GetSlack() and GetOptFreq() can be highly non-trivial. Specifically, GetSlack() calculates slack of a task off the CP from the difference between the latest and the earliest end time of the task. GetOptFreq() calculates the optimal ideal frequency to eliminate slack of a task from the mapping between frequency and execution time for each type of tasks.

4.2. Scheduled Communication Approach

One effective and straightforward solution to save energy for task-parallel applications is to set CPU frequency to high during computation, while set it to low during communication, given the fact that large-message communication is not bound by CPU performance while the computation is, so the peak

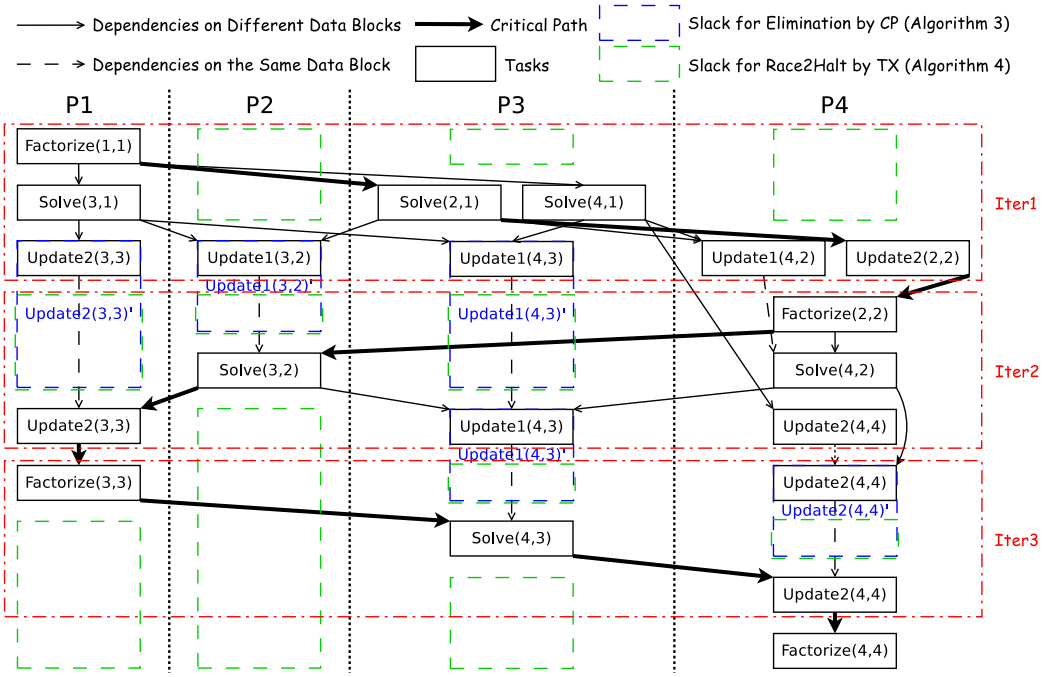

Figure 4: DAG Representation of Task and Slack Scheduling of CP and TX Approaches for the 4 × 4 Blocked Cholesky Factorization in Figure 3 on a 2 × 2 Process Grid Using 2-D Block Cyclic Data Distribution.

CPU performance is not necessary during the communication. Although substantial energy savings can be achieved from this Scheduled Communication (*SC*) approach [25] [8], it leaves potential energy saving opportunities from other types of slack (e.g., see slack shown in Figure 4 on page 15) untapped. More energy savings can be fulfilled via fine-grained analysis of execution characteristics of HPC applications, in particular during non-communication. Next we present two well-designed approaches that take advantage of computation slack to further gain energy savings. Note since the *SC* approach does not conflict with solutions exploiting slack from non-communication, it thus can be incorporated with the next two solutions seamlessly to maximize energy savings.

4.3. Critical Path Approach vs. TX Approach

Given a detected CP (e.g., via static analysis [3] or local information analysis [9]) for task-parallel applications, the Critical Path (*CP*) approach saves energy as shown in Algorithm 3: For all tasks on the CP, the CPU

Algorithm 3 *DVFS Scheduling Algorithm Using CP*

```
DVFS_CP(CritPath, task, FreqSet)
 1: if (task ∈ CritPath || TDS_out(task) != ∅) then
 2:    SetFreq(f_h)
 3: else
 4:    slack ← GetSlack(task)
 5:    if (slack > 0) then
 6:       f_opt ← GetOptFreq(task, slack)
 7:    if (f_l ≤ f_opt ≤ f_h) then
 8:       if (f_opt ∉ FreqSet) then
 9:          SetFreq(⌊f_opt⌋, ⌈f_opt⌉, ratio)
10:       else SetFreq(f_opt)
11:    else if (f_opt < f_l) then
12:       SetFreq(f_l)
13: end if
```

operating frequency is set to the highest for attaining the peak CPU performance, while for tasks not on the CP with the total slack larger than zero (e.g., tasks with no outgoing *explicit* data dependencies in Figure 4), lowering frequency appropriately is performed to dilate the tasks into their slack as much as possible, without incurring performance loss of the applications. Due to the discrete domain of available CPU frequencies defined for DVFS, if the calculated optimal frequency that can eliminate slack lies in between two available neighboring frequencies, the two frequencies can be employed to approximate it by calculating a ratio of durations operating at the two frequencies. The two frequencies are then assigned to the durations separately based on the ratio. Lines 7-9 in Algorithm 3 sketch the frequency approximation method [31] [9]. The ratio of split frequencies is calculated via prior knowledge of the mapping between frequency and execution time of different types of tasks. Note that we denote the two neighboring available frequencies of f_{opt} as $\lfloor f_{opt} \rfloor$ and $\lceil f_{opt} \rceil$.

Different from the *CP* approach that reduces CPU frequency for tasks off the CP to eliminate slack for saving energy without performance loss, `TX` employs the *race-to-halt* mechanism that leverages two TDS of each task (TDS_{in} and TDS_{out}) to determine the timing of *race* and *halt*, as shown in Algorithm 4. Respecting data dependencies, one *dependent* task cannot start

Algorithm 4 *DVFS Scheduling Algorithm Using TX*

```
DVFS_TX(task, CurFreq)
 1: while (TDS_in(task) != ∅) do
 2:    if (CurFreq != f_l) then
 3:       SetFreq(f_l)
 4:    if (Recv(DoneFlag, t_1)) then
 5:       delete(TDS_in(task), t_1)
 6: end while
 7: SetFreq(f_h)
 8: if (IsFinished(task)) then
 9:    foreach t_2 ∈ TDS_out(task) do
10:       Send(DoneFlag, t_2)
11:    SetFreq(f_l)
12: end if
```

until the finish of its *depended* task. `TX` keeps the *dependent* task staying at the lowest frequency, i.e., *halt*, until all its *depended* tasks have finished when it may start, and then allows the *dependent* task to work at the highest frequency to complete as soon as possible, i.e., *race*, before being switched back to the low-power state. Upon completion, a task sends a $DoneFlag$ to all its *dependent* tasks to notify them that data needed has been processed and ready for use. A *dependent* task is retained at the lowest frequency while waiting for $DoneFlags$ from all its *depended* tasks, and removes the dependency to a *depended* task from its TDS_{in}, once a $DoneFlag$ from the *depended* task is received.

Defective Prediction Mechanism. Although effective, the OS level *CP* approach essentially depends on the workload prediction mechanism: Execution characteristics in the upcoming interval can be predicted using prior execution information, e.g., execution traces in recent intervals. However, this prediction mechanism may not necessarily be reliable and lightweight: (a) For applications with variable execution characteristics, such as dense matrix factorizations. Execution time of iterations of the core loop shrinks as the remaining unfinished matrices become smaller. Consequently dynamic prediction on execution characteristics such as task runtime and workload distribution can be inaccurate, which thus leads to error-prone slack estimation; (b) for applications with random execution patterns, such as applica-

tions relying on random numbers that could lead to variation of control flow at runtime, which can be difficult to capture. Since the predictor needs to determine reproducible execution patterns at the beginning of one execution, it can be costly for obtaining accurate prediction results in both cases above. Given the fact that no energy savings can be fulfilled until the prediction phase is finished, considerable potential energy savings may be wasted for accurate but lengthy prediction as such at OS level.

There exist numerous studies on history-based workload prediction, which can be generally considered as two variants: the simplest and mostly commonly-used prediction algorithm PAST [32] and its enhanced algorithms [33] [34] [20] [11], where PAST works as follows:

$$W'_{i+1} = W_i \tag{1}$$

where W'_{i+1} is the next executed workload to predict, and W_i is the current measured workload. For applications with stable or slowly-varying execution patterns, the straightforward PAST algorithm can work well with little performance loss and high accuracy. It is however not appropriate for handling a considerable amount of variation in execution patterns. Many enhanced algorithms have been proposed to produce more accurate workload prediction for such applications. For instance, the RELAX algorithm employed in CPU MISER [11] exploits both prior predicted profiles and current runtime measured profiles as follows:

$$W'_{i+1} = (1 - \lambda)W'_i + \lambda W_i \tag{2}$$

where λ is a relaxation factor for adjusting the extent of dependent information of the current measurement. This enhanced prediction mechanism can also be error-prone and costly for parallel Cholesky, LU, and QR factorizations due to the use of 2-D block cyclic data distribution. As shown in Figure 4, across loop iterations (highlighted by red dashed boxes) different types of tasks can be distributed to a process. For instance, in the first iteration, $\mathsf{P_1}$ is assigned three tasks Factorize(1,1), Solve(3,1), and Update2(3,3), while in the second iteration, it is only assigned one task Update2(3,3). Although empirically for a large global matrix, tasks are distributed to one specific process alternatingly (e.g., Factorize(i,i) ($\mathsf{i} = 1, 3, 5, \ldots, 2\mathsf{s}+1, \mathsf{s} \geq 0$) can be all distributed to P1), the RELAX algorithm needs to adjust the value of λ alternatingly as well, which can be very costly and thus diminish potential energy savings. Length variation of iterations due to the shrinking remaining unfinished matrices further brings complexity to workload prediction.

TX successfully circumvents the shortcomings from the OS level workload

prediction, by leveraging the TDS-based *race-to-halt* strategy. `TX` essentially differs from the *CP* approach since it allows tasks to complete as soon as possible before entering the low-power state, and thus does not require any workload prediction mechanism. `TX` works at library level and thus benefits from: (a) obtaining application-specific knowledge easily that can be utilized to determine accurate timing for DVFS, and (b) restricting the scope of source modification and recompilation to library level, compared to application level solutions.

Potential Energy Savings from Load Imbalance. Due to the existence of load imbalance regardless of load balancing techniques applied, and also data dependencies among inter-process parallel tasks, not all tasks can start to work and finish at the same time, as shown in Figure 4. At OS level, the *SC* approach and the *CP* approach only work when tasks are being executed, which leaves potential energy savings untapped during the time otherwise. Specifically, the *SC* approach switches frequency to high and low at the time when computation and communication tasks start respectively, and do nothing during the time other than computation and communication. Likewise, the *CP* approach assigns appropriate frequencies for tasks on/off the CP individually. Therefore due to its nature of keeping the peak CPU performance for the tasks on the CP and dilating the tasks off the CP into its slack via frequency reduction, switching frequency is not feasible when for one process no tasks have started or all tasks have already finished.

Due to its nature of *race-to-halt*, `TX` can start to save energy even before any tasks in a process are executed, and after all tasks in a process have finished while there exist unfinished tasks in other processes. In Figure 4, the durations only covered by green dashed boxes highlights the additional energy savings fulfilled by `TX`, where due to data dependencies, a task cannot start yet while waiting for its depended task to finish first, or all tasks in one process have already finished while some tasks in other processes are still running.

Energy Saving Capability Analysis. Next we formally prove that compared to the classic *CP* approach, `TX` is comparable to it in energy saving capability. Given the following two energy saving strategies, towards a task t with an execution time T and slack T' at the peak CPU performance, we calculate the total nodal system energy consumption for both strategies, i.e., $E(\mathsf{S}_1)$ and $E(\mathsf{S}_2)$ formally below:

- Strategy I (Race-to-halt): Execute t at the highest frequency f_h until

Table 2: Notation in Energy Saving Analysis.

E	The total nodal energy consumption of all components
P	The total nodal power consumption of all components
$P_{dynamic}$	Dynamic power consumption in the running state
$P_{leakage}$	Static/leakage power consumption in any states
T	Execution time of a task at the peak CPU performance
T'	Slack of executing a task at the peak CPU performance
A	Percentage of active gates in a CMOS-based chip
C	The total capacitive load in a CMOS-based chip
f	Current CPU operating frequency
V	Current CPU supply voltage
V'	Supply voltage of components other than CPU
I_{sub}	CPU subthreshold leakage current
I'_{sub}	non-CPU component subthreshold leakage current
f_m	Available optimal frequency assumed to eliminate T'
V_h	The highest supply voltage corresponding to f_h set by DVFS
V_l	The lowest supply voltage corresponding to f_l set by DVFS
V_m	Supply voltage corresponding to f_m set by DVFS
n	Ratio between original runtime and slack of a task

the finish of t and then switch to the lowest frequency f_l, i.e., run in T at f_h and then run in T' at f_l;

- Strategy II (CP-aware): Execute t at the optimal frequency f_m with which T' is eliminated, i.e., run in $T + T'$ at f_m (without loss of generality, assume T' can be eliminated using an available frequency f_m without frequency approximation).

For simplicity, let us assume the tasks for the use of DVFS are compute-intensive (memory-intensive tasks can be discussed with minor changes in the model), i.e., $T + T' = nT$, when $f_m = \frac{1}{n} f_h$, where $1 \leq n \leq \frac{f_h}{f_l}$. Considering the nodal power consumption P, we model it formally as follows:

$$P = P^{CPU}_{dynamic} + P^{CPU}_{leakage} + P^{other}_{leakage} \tag{3}$$

$$P_{dynamic} = ACfV^2 \tag{4}$$

$$P_{leakage} = I_{sub}V \tag{5}$$

Then, substituting Equations 4 and 5 into Equation 3 yields:

$$P = ACfV^2 + I_{sub}V + I'_{sub}V' \tag{6}$$

In our scenario, $P^{other}_{leakage} = I'_{sub}V'$ is independent of CPU frequency and voltage scaling, and thus can be regarded as a constant in Equation 6, so we denote $P^{other}_{leakage}$ as P_c for short. Further, although subthreshold leakage current I_{sub} has an exponential relationship with threshold voltage, results presented in [35] indicate that I_{sub} converges to a constant after a certain threshold voltage value. Without loss of generality, we treat $P^{CPU}_{leakage} = I_{sub}V$ as a function of supply voltage V only. Thus, we model the nodal energy consumption E_{node} for both strategies individually below:

$$\begin{aligned} E(\mathsf{S_1}) &= \overline{P(\mathsf{S_1})} \times T + \overline{P'(\mathsf{S_1})} \times T' \\ &= (ACf_hV_h^2 + I_{sub}V_h + P_c)T + (ACf_lV_l^2 + I_{sub}V_l + P_c)T' \\ &= AC(f_hV_h^2T + f_lV_l^2T') + I_{sub}(V_hT + V_lT') + P_c(T + T') \end{aligned} \tag{7}$$

$$\begin{aligned} E(\mathsf{S_2}) &= \overline{P(\mathsf{S_2})} \times (T + T') \\ &= (ACf_mV_m^2 + I_{sub}V_m + P_c)(T + T') \\ &= ACf_mV_m^2(T + T') + I_{sub}V_m(T + T') + P_c(T + T') \end{aligned} \tag{8}$$

We obtain the difference between energy costs of both strategies by dividing Equation 8 by Equation 7:

$$\frac{E(\mathsf{S_2})}{E(\mathsf{S_1})} = \frac{ACf_mV_m^2(T + T') + I_{sub}V_m(T + T') + P_c(T + T')}{AC(f_hV_h^2T + f_lV_l^2T') + I_{sub}(V_hT + V_lT') + P_c(T + T')} \tag{9}$$

Substituting the assumption that $T' = (n-1)T$ and $f_m = \frac{1}{n}f_h$ into both numerator and denominator yields the following simplified formula:

$$\begin{aligned} \frac{E(\mathsf{S_2})}{E(\mathsf{S_1})} &= \frac{AC\frac{1}{n}f_hV_m^2\left(1 + (n-1)\right) + I_{sub}V_m\left(1 + (n-1)\right) + P_c\left(1 + (n-1)\right)}{AC\left(f_hV_h^2 + f_lV_l^2\left(n-1\right)\right) + I_{sub}\left(V_h + V_l\left(n-1\right)\right) + P_c\left(1 + (n-1)\right)} \\ &= \frac{ACf_hV_m^2 + nI_{sub}V_m + nP_c}{AC\left(f_hV_h^2 + f_lV_l^2\left(n-1\right)\right) + I_{sub}\left(V_h + V_l\left(n-1\right)\right) + nP_c} \end{aligned} \tag{10}$$

In Equation 10, the denominator is a function of the variable n only. It is clear that it is a monotonically increasing function for n, whose minimum value is attained when $n = 1$, i.e., when slack T' equals 0. Given the fact that supply voltage has a positive correlation with (not strictly proportional to)

Table 3: Frequency-Voltage Pairs for Different Processors (Unit: Frequency (GHz), Voltage (V)).

Gear	AMD Opteron 2380		AMD Opteron 846 and AMD Athlon64 3200+		AMD Opteron 2218		Intel Pentium M		Intel Core i7-2760QM	
	Freq.	Volt.	Freq.	Volt.	Freq.	Volt.	Freq.	Volt.	Freq.	Volt.
0	2.5	1.300	2.0	1.500	2.4	1.250	1.4	1.484	2.4	1.060
1	1.8	1.200	1.8	1.400	2.2	1.200	1.2	1.436	2.0	0.970
2	1.3	1.100	1.6	1.300	1.8	1.150	1.0	1.308	1.6	0.890
3	0.8	1.025	0.8	0.900	1.0	1.100	0.8	1.180	0.8	0.760

operating frequency, scaling up/down frequency results in voltage up/down accordingly as shown in Table 3. Therefore for the numerator of Equation 10, the greater n is, the smaller f_m and V_m are, provided $f_m = \frac{1}{n} f_h$ and the above fact. It is thus complicated to determine the monotonicity of the numerator. As a matter of fact, state-of-the-art CMOS technologies allow insignificant variation of voltage as frequency scales (see Table 3). Consequently the term $ACf_hV_m^2$ within the numerator does not decrease much together with the increase of n. Moreover, the ratio between the highest and the lowest frequencies determines the upper bound of n ($1 \leq n \leq \frac{f_h}{f_l}$), so the other two terms within the numerator cannot increase significantly as n goes up.

Example. From the operating points of various processors shown in Table 3, we can calculate numerical energy savings for different values of n and a specific processor configuration that quantify the energy efficiency of both energy saving strategies. For the AMD Opteron 2218 processor, given a task with the execution time T and slack $0.25T$, i.e., $n = 1.25$, for eliminating the slack, 1.8 GHz is adopted as the operating frequency for running the task, and thus the numerator of Equation 10 equals $AC\,(2.4 \times 1.25^2 + (1.25 - 1) \times 1.0 \times 1.1^2) + I_{sub}\,(1.25 + (1.25 - 1) \times 1.1) + 1.25P_c = 4.0525AC + 1.525I_{sub} + 1.25P_c$, while the denominator of Equation 10 equals $AC \times 2.4 \times 1.15^2 + 1.25 \times 1.15I_{sub} + 1.25P_c = 3.174AC + 1.4375I_{sub} + 1.25P_c$. We can see that the coefficients of the corresponding term P_c between the numerator and the denominator are the same, and the coefficients of both corresponding terms AC and I_{sub} between the numerator and the denominator do not differ much. Therefore the result of $\frac{E(\mathsf{S}_2)}{E(\mathsf{S}_1)}$ would be a value close to 1, which means Strategy I is comparable to Strategy II in energy efficiency, regardless of the defective prediction mechanism.

5. Implementation and Evaluation

We have implemented `TX`, and for comparison purposes, the library level *SC* approach to evaluate the effectiveness of `TX` to save energy during non-communication slack. For comparing with the OS level *SC* and *CP* approaches, we communicated with the authors of Adagio [9] and Fermata [8] and received the latest version of the implementation of both. We also compare with another OS level solution CPUSpeed [36], an interval-based DVFS scheduler that scales CPU performance according to runtime CPU utilization during the past interval. Regarding future workload prediction, essentially Adagio and Fermata leverage the PAST algorithm [32], and CPUSpeed uses a prediction algorithm similar to the RELAX algorithm employed in CPU MISER [11]. With application-specific knowledge known, all library level solutions do not need the workload prediction mechanism. In the later discussion, we denote the above approaches as follows:

- `Orig`: The original runs of different-scale parallel Cholesky, LU, and QR factorizations without any energy saving approaches;
- `SC_lib`: The library level implementation of the *SC* approach;
- `Fermata`: The OS level implementation of the *SC* approach based on the PAST workload prediction algorithm;
- `Adagio`: The OS level implementation of the *CP* approach based on the PAST workload prediction algorithm, where `Fermata` is incorporated;
- `CPUSpeed`: The OS level implementation of the *SC* approach based on a workload prediction algorithm similar to the RELAX algorithm;
- `TX`: The library level implementation of the *race-to-halt* approach based on TDS analysis, where `SC_lib` is incorporated.

The goals of the evaluation are to demonstrate that: (a) `TX` is able to save energy effectively and efficiently for applications with variable execution characteristics such as parallel Cholesky, LU, and QR factorizations, while OS level prediction-based solutions cannot maximize energy savings, and (b) `TX` only incurs negligible performance loss, similar as the compared OS level solutions. We did not compare with application level solutions, since they essentially fulfill the same energy efficiency as library level solutions, with source

modification and recompilation at application level. Working at library level, `TX` was deployed in a distributed manner to each core/process. As the additional functionality of saving energy, the implementation of `TX` was embedded into a rewritten version of ScaLAPACK [21], a widely used high performance and scalable numerical linear algebra library for distributed-memory architectures. In particular, library level source modification and recompilation were conducted to the pdpotrf(), pdgetrf(), and pdgeqrf() routines, which perform parallel Cholesky, LU, and QR factorizations, respectively.

Table 4: Hardware Configuration for All Experiments.

Cluster	HPCL	ARC
System Size (# of Nodes)	8	108
Processor	2×Quad-core AMD Opteron 2380	2×8-core AMD Opteron 6128
CPU Freq.	0.8, 1.3, 1.8, 2.5 GHz	0.8, 1.0, 1.2, 1.5, 2.0 GHz
Memory	8 GB RAM	32 GB RAM
Cache	128 KB L1, 512 KB L2, 6 MB L3	128 KB L1, 512 KB L2, 12 MB L3
Network	1 GB/s Ethernet	40 GB/s InfiniBand
OS	CentOS 6.2, 64-bit Linux kernel 2.6.32	CentOS 5.7, 64-bit Linux kernel 2.6.32
Power Meter	PowerPack	Watts up? PRO

5.1. Experimental Setup

We applied all five energy efficient approaches to the parallel Cholesky, LU, and QR factorizations with five different global matrix sizes each to assess our goals. Experiments were performed on two power-aware clusters: HPCL and ARC. Table 4 lists the hardware configuration of the two clusters. Note that we measured the total dynamic and leakage energy consumption of distributed runs using PowerPack [18], a comprehensive software and hardware framework for energy profiling and analysis of high performance systems and applications. The total of static and dynamic power consumption was measured using Watts up? PRO [37]. Both energy and power consumption are the total energy and power costs respectively on all involved components of one compute node, such as CPU, memory, disk, motherboard, etc. Since each set of three nodes of the ARC cluster share one power meter, power consumption measured is for the total power consumption of three nodes, while energy consumption measured is for all energy costs collected from all eight nodes of the HPCL cluster. CPU DVFS was implemented via the

CPUFreq infrastructure [16] that directly modifies CPU frequency system configuration files.

5.2. Results

In this section, we present experimental results on power, energy, and performance efficiency and trade-offs individually by comparing `TX` with the other energy saving approaches.

Power Savings. First we evaluate the capability of power savings from the five energy saving approaches for parallel Cholesky, LU, and QR factorizations on the ARC cluster (due to the similarity of results, data for parallel LU and QR factorizations is not shown), where the power consumption was measured by sampling at a constant rate through the execution of the applications. Figure 5 depicts the total power consumption of three nodes (out of sixteen nodes in use) running parallel Cholesky factorization with different approaches using a 160000 $\times$ 160000 global matrix, where we select time durations of the first few iterations. Note that parallel Cholesky factorization (the core loop) performs alternating computation and communication with decreasing execution time of each iteration, as the remaining unfinished matrix shrinks. Thus we can see that for all approaches, from left to right, the durations of computation (the peak power values) decrease as the factorization proceeds.

Among the seven runs (including the theoretical one `CP_theo`), there exist four power variation patterns: (a) `Orig` and `CPUSpeed` – employed the same highest frequency for both computation and communication, resulting almost constant power consumption around 950 Watts; (b) `SC_lib`, `Fermata`, and `Adagio` – lowered down frequency during the communication, i.e., the five low-power durations around 700 Watts, and resumed the peak CPU performance during the computation; (c) `CP_theo` – not only scheduled low power states for communication, but also slowed down computation to eliminate computation slack – this is a theoretical value curve instead of real measurement, which is how OS level approaches such as `Adagio` is supposed to save more power as a *CP-aware* approach based on accurate workload prediction; and (d) `TX` – employed the *race-to-halt* strategy to lower down CPU performance for all durations other than computation.

Specifically, upon the prediction algorithm that inspects dynamic prior CPU utilization periodically for workload prediction, `CPUSpeed` failed to produce accurate prediction and scale CPU power states accordingly: It kept the peak CPU performance all the time. Either relying on knowing application

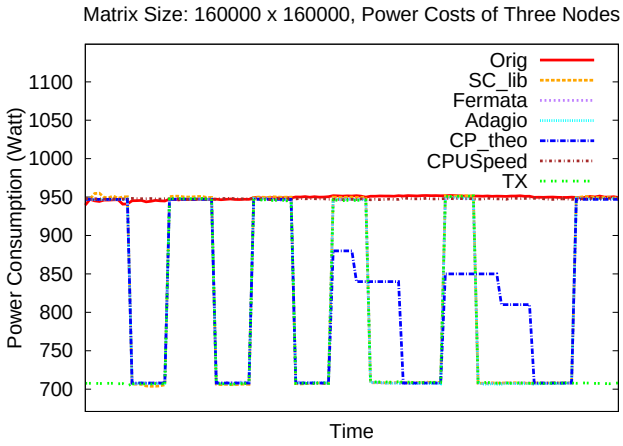

Figure 5: Power Consumption of Parallel Cholesky Factorization with Different Energy Saving Approaches on the ARC Cluster using 16 × 16 Process Grid.

characteristics (`SC_lib`) or detecting MPI communication calls (`Fermata` and `Adagio`), all three approaches can identify durations of communication and apply DVFS decisions accordingly. As discussed earlier, solutions that only slow down CPU during communication are semi-optimal. `Adagio` and `TX` are expected to utilize computation slack for achieving additional energy savings. Due to the defective OS level prediction mechanism, `Adagio` failed to predict behavior of future tasks and calculate computation slack accurately. Consequently no low-power states were switched to during computation for `Adagio`. In contrast, we provide a theoretical value curve `CP_theo` that calculates computation slack effectively, and lower power states were switched to eliminate the slack, i.e., the four medium-power durations around 850 Watts during the third and the fourth computation as the blue line shows. Different from the solutions that save energy via slack reclamation, `TX` relies on the *race-to-halt* mechanism where computation is conducted at the peak CPU performance and the lowest CPU frequency is employed immediately after the computation. Therefore during computation slack, we can observe low-power states were switched to by `TX`. Moreover, the nature of *race-to-halt* also guarantees no high-power states are employed during the waiting durations resulting from data dependencies and load imbalance, i.e., the two low-power durations in green where the application starts and ends. This indicates that `TX` is able to gain additional energy savings that all other approaches cannot exploit: Processes have to stay at the high-power state at

the beginning/ending of the execution.

Energy Savings. Next we compare energy savings achieved by all five approaches (not including `CP_theo`) on parallel Cholesky, LU, and QR factorizations on the HPCL cluster as shown in Figure 6, where the energy consumption was measured by recording on/off the collection of power and time costs when an application starts/ends. For eliminating errors from scalability, we collected energy and time data of five matrix factorizations with different global matrix sizes ranging from 5120 to 25600, respectively. Considerable energy savings are achieved by all approaches except for `CPUSpeed` on all Cholesky, LU, and QR with similar energy saving trend: `TX` prevails over all other approaches with higher energy efficiency, while `SC_lib`, `Fermata`, and `Adagio` have similar energy efficiency. Overall, for Cholesky, `TX` can save energy 30.2% on average and up to 33.8%; for LU and QR, `TX` can achieve 16.0% and 20.0% on average and up to 20.4% and 23.4% energy savings, respectively. Due to the reasons discussed for power savings, `Adagio` only achieves similar energy savings as `SC_lib` and `Fermata`, without fulfilling additional energy savings from slack reclamation of computation. With application-specific knowledge instead of workload prediction, `TX` manages to achieve energy savings during both computation and communication slack. Moreover, `TX` benefits from the advantage of saving additional energy during possible load imbalance other approaches cannot exploit. Next we further evaluate such energy savings by increasing load imbalance in the applications.

Effects of Block Size. As discussed earlier, the additional energy savings can be achieved from load imbalance, i.e., the durations only covered by green dashed boxes as shown in Figure 4. Empirically, regardless of the workload partition techniques employed, load imbalance can grow due to larger tasks, longer communication, etc. For manifesting the strength of `TX` in achieving additional energy savings for *completeness*, we deliberately imbalance the workload through expanding tasks by using greater block sizes for Cholesky, while keeping the default block size for LU and QR. As shown in Figure 6, the average energy savings fulfilled by `TX` for Cholesky (30.2%) are consequently greater than LU and QR (16.0% and 20.0%). Compared to the second most effective approach `Adagio`, `TX` can save Cholesky 12.8% more energy on average.

Performance Loss. Figure 6 also illustrates performance loss from different energy saving approaches against the original runs. We can see `TX` only incurs negligible time overhead: 3.8%, 3.1%, 3.7% on average for Cholesky,

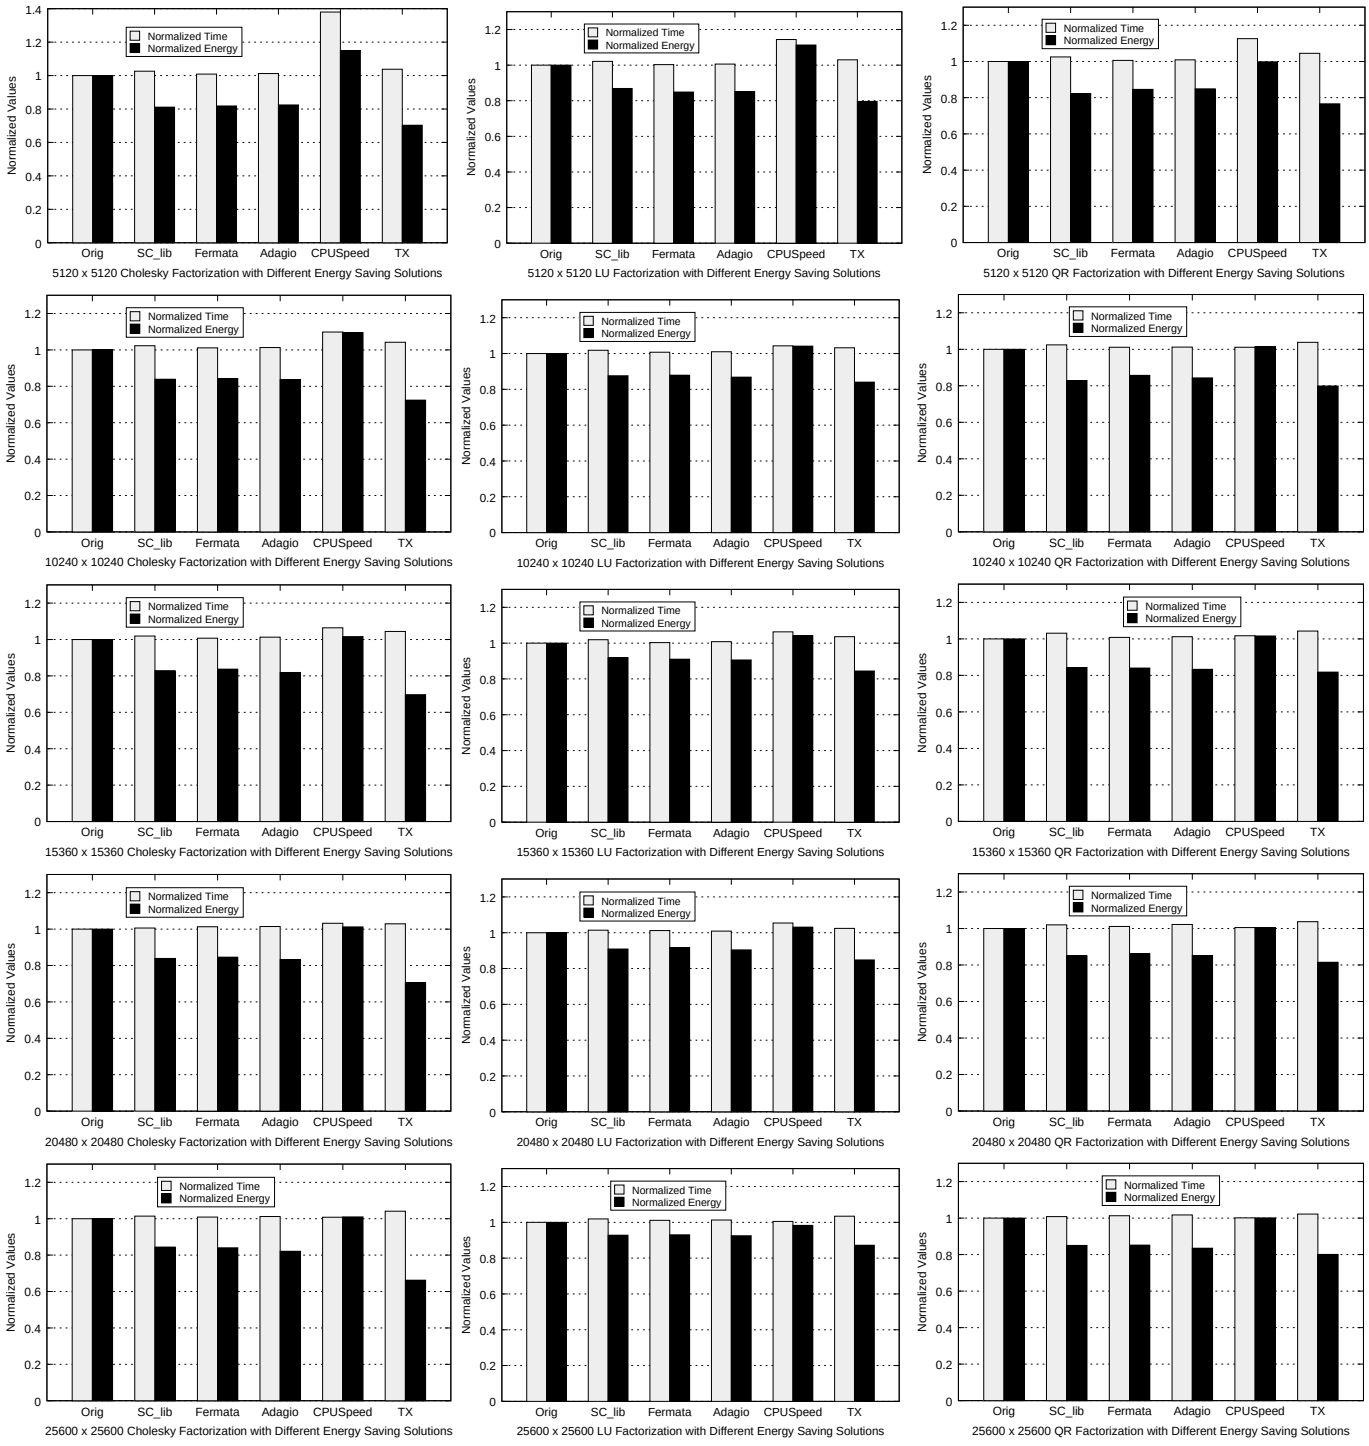
Normalized Time
Normalized Energy
Normalized Values
Orig
SC_lib
Fermata
Adagio
CPUSpeed
TX
5120 x 5120 Cholesky Factorization with Different Energy Saving Solutions
5120 x 5120 LU Factorization with Different Energy Saving Solutions
5120 x 5120 QR Factorization with Different Energy Saving Solutions
10240 x 10240 Cholesky Factorization with Different Energy Saving Solutions
10240 x 10240 LU Factorization with Different Energy Saving Solutions
10240 x 10240 QR Factorization with Different Energy Saving Solutions
15360 x 15360 Cholesky Factorization with Different Energy Saving Solutions
15360 x 15360 LU Factorization with Different Energy Saving Solutions
15360 x 15360 QR Factorization with Different Energy Saving Solutions
20480 x 20480 Cholesky Factorization with Different Energy Saving Solutions
20480 x 20480 LU Factorization with Different Energy Saving Solutions
20480 x 20480 QR Factorization with Different Energy Saving Solutions
25600 x 25600 Cholesky Factorization with Different Energy Saving Solutions
25600 x 25600 LU Factorization with Different Energy Saving Solutions
25600 x 25600 QR Factorization with Different Energy Saving Solutions

Figure 6: Energy and Performance Efficiency of Parallel Cholesky, LU, and QR Factorizations on the HPCL Cluster with Different Global Matrix Sizes and Energy Saving Approaches using 8×8 Process Grid.

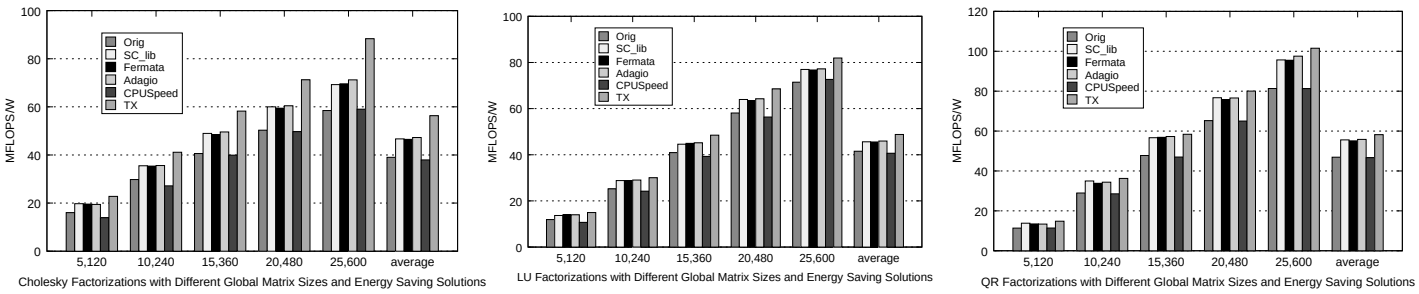

Figure 7: Energy and Performance Trade-off of Parallel Cholesky, LU, and QR Factorizations on the HPCL Cluster with Different Global Matrix Sizes and Energy Saving Approaches using 8×8 Process Grid.

LU, and QR individually, similar to the time overhead of all other approaches except for `CPUSpeed`. The minor performance loss on employing these solutions is primarily originated from three aspects: (a) Although large-message communication is not CPU-bound, pre-computation required for starting up a communication link before any data transmission is necessary and is affected by CPU performance, so the low-power state during communication can slightly degrade performance; (b) switching CPU frequency via DVFS is essentially implemented by modifying CPU frequency system configuration files, and thus slight time overhead is incurred from the in-memory file read/write operations [14]; and (c) CPU frequency transition latency is required for the newly-set frequency to take effect. Further, `TX` suffers from minor performance loss from TDS analysis, including TDS generation and maintaining TDS for each task. The high time overhead of `CPUSpeed` is another reason for its little and even negative energy savings besides the defective prediction mechanism at OS level.

Energy/Performance Trade-off. An optimal energy saving approach requires to achieve the maximal energy savings with the minimal performance loss. Per this requirement, energy-performance integrated metrics are widely employed to quantify if the energy efficiency achieved and the performance loss incurred meanwhile are well-balanced. We adopt Energy-Delay Product (EDP) to evaluate the overall energy and performance trade-off of the five approaches, in terms of MFLOPS/W, which equals the amount of floating-point operations per second within the unit of one Watt, i.e., the greater value it is, the better efficiency is fulfilled. As shown in Figure 7, compared to other approaches, `TX` is able to fulfill the most balanced trade-off between the energy and performance efficiency achieved. Specifically, `TX` has

higher MFLOPS/W values for Cholesky compared to LU and QR, due to the higher energy savings achieved from the more imbalanced workload in Cholesky without additional performance loss.

6. Related Work

Numerous other types of energy efficient DVFS scheduling algorithms have been proposed, but only a few of them were designed for high performance scientific computing. We next detail them in the following categories *OS-level*, *Application-level*, *Simulation-based*.

OS-level. There exist a large body of OS level energy efficient approaches for high performance scientific applications. Lim *et al.* [25] developed a runtime system that dynamically and transparently reduces CPU power for communication phases to minimize energy-delay product. Ge *et al.* [11] proposed a runtime system and an integrated performance model for achieving energy efficiency and constraining performance loss through performance modeling and prediction. Rountree *et al.* [8] developed a *SC* approach that employs a linear programming solver collecting communication trace and power characteristics for generating an energy saving scheduling. Subsequent work [9] presented another runtime system by improving and extending previous classic scheduling algorithms and achieved significant energy savings with extremely limited performance loss.

Application-level. Kappiah *et al.* [6] introduced a scheduled iteration method that computes the total slack per processor per timestep, then scheduling CPU frequency for the upcoming timestep. Liu *et al.* [38] presented a technique that tracks the idle durations for one processor to wait for others to reach the same program point, and utilizes this information to reduce the idle time via DVFS without performance loss. Tan *et al.* [13] proposed an adaptively aggressive scheduling strategy for data intensive applications with moderated performance trade-off using speculation. Subsequent work [14] proposed an adaptive memory-aware strategy for distributed matrix multiplication that trades grouped computation/communication with memory costs for less overhead on employing DVFS. Liu *et al.* [39] proposed a power-aware static mapping technique to assign applications for a CPU/GPU heterogeneous system that reduces power and energy costs via DVFS on both CPU and GPU, with timing requirements satisfied.

Simulation-based. There exist some efforts on improving energy efficiency

for numerical linear algebra operations like Cholesky/LU/QR factorization, but most of them either are based on simulation or only work for a single multicore machine. Few studies have been conducted on power/energy efficient matrix factorizations running on distributed-memory architectures. Slack reclamation methods such as Slack Reduction and Race-to-Idle algorithms [27] [26] have been proposed to save energy for dense linear algebra operations on shared-memory multicore processors. Instead of running benchmarks on real machines, a power-aware simulator, in charge of runtime scheduling to achieve task level parallelism, was employed to evaluate the proposed power-control policies for linear algebra operations. DVFS techniques used in their approaches were also simulated. Subsequent work [40] leveraged DVFS to optimize task-parallel execution of a collection of dense linear algebra tasks on shared-memory multicore architectures, where experiments were performed at thread level. Since no communication was involved, these approaches did not achieve significant energy savings due to no utilization of slack from communication latency.

7. Conclusions

The looming overloaded energy consumption of high performance scientific computing brings significant challenges to green computing in this era of ever-growing power costs for large-scale HPC systems. DVFS techniques have been widely employed to improve energy efficiency for task-parallel applications. With high generality, OS level solutions are regarded as feasible energy saving approaches for such applications. We observe for applications with variable execution characteristics such as parallel Cholesky, LU, and QR factorizations, OS level solutions suffer from the defective prediction mechanism and untapped potential energy savings from possible load imbalance, and thus cannot optimize the energy efficiency. Giving up partial generality, the proposed library level approach `TX` is evaluated to save more energy with little performance loss for parallel Cholesky, LU, and QR factorizations on two power-aware clusters compared to classic OS level solutions.

References

[1] TOP500 Supercomputer Lists, http://www.top500.org/.

[2] Green500 Supercomputer Lists, http://www.green500.org/.

[3] G. Chen, K. Malkowski, M. Kandemir, P. Raghavan, Reducing power with performance constraints for parallel sparse applications, in: IPDPS, 2005, pp. 1–8.

[4] V. W. Freeh, D. K. Lowenthal, Using multiple energy gears in MPI programs on a power-scalable cluster, in: PPoPP, 2005, pp. 164–173.

[5] R. Ge, X. Feng, K. W. Cameron, Performance-constrained distributed DVS scheduling for scientific applications on power-aware clusters, in: SC, 2005, p. 34.

[6] N. Kappiah, V. W. Freeh, D. K. Lowenthal, Just in time dynamic voltage scaling: Exploiting inter-node slack to save energy in MPI programs, in: SC, 2005, p. 33.

[7] R. Springer, D. K. Lowenthal, B. Rountree, V. W. Freeh, Minimizing execution time in MPI programs on an energy-constrained, power-scalable cluster, in: PPoPP, 2006, pp. 230–238.

[8] B. Rountree, D. K. Lowenthal, S. Funk, V. W. Freeh, B. R. de Supinski, M. Schulz, Bounding energy consumption in large-scale MPI programs, in: SC, 2007, pp. 1–9.

[9] B. Rountree, D. K. Lowenthal, B. R. de Supinski, M. Schulz, V. W. Freeh, T. Bletsch, Adagio: Making DVS practical for complex HPC applications, in: ICS, 2009, pp. 460–469.

[10] K. H. Kim, R. Buyya, J. Kim, Power aware scheduling of bag-of-tasks applications with deadline constraints on DVS-enabled clusters, in: CCGrid, 2007, pp. 541–548.

[11] R. Ge, X. Feng, W.-C. Feng, K. W. Cameron, CPU MISER: A performance-directed, run-time system for power-aware clusters, in: ICPP, 2007, p. 18.

[12] X. Wang, X. Fu, X. Liu, Z. Gu, Power-aware CPU utilization control for distributed real-time systems, in: RTAS, 2009, pp. 233–242.

[13] L. Tan, Z. Chen, Z. Zong, R. Ge, D. Li, A2E: Adaptively aggressive energy efficient DVFS scheduling for data intensive applications, in: IPCCC, 2013, pp. 1–10.

[14] L. Tan, L. Chen, Z. Chen, Z. Zong, R. Ge, D. Li, HP-DAEMON: High performance distributed adaptive energy-efficient matrix-multiplication, in: ICCS, 2014, pp. 599–613.

[15] H. David, C. Fallin, E. Gorbatov, U. R. Hanebutte, O. Mutlu, Memory power management via dynamic voltage/frequency scaling, in: ICAC, 2011, pp. 31–40.

[16] CPUFreq - CPU Frequency Scaling, https://wiki.archlinux.org/index.php/CPU_Frequency_Scaling.

[17] NVIDIA System Management Interface (nvidia-smi), https://developer.nvidia.com/nvidia-system-management-interface/.

[18] R. Ge, X. Feng, S. Song, H.-C. Chang, D. Li, K. W. Cameron, PowerPack: Energy profiling and analysis of high-performance systems and applications, IEEE Transactions on Parallel and Distributed Systems 21 (5) (2010) 658–671.

[19] A. Miyoshi, C. Lefurgy, E. V. Hensbergen, R. Rajamony, R. Rajkumar, Critical power slope: Understanding the runtime effects of frequency scaling, in: ICS, 2002, pp. 35–44.

[20] C.-H. Hsu, W.-C. Feng, A power-aware run-time system for high-performance computing, in: SC, 2005, p. 1.

[21] Scalable Linear Algebra PACKage, http://www.netlib.org/scalapack/.

[22] DPLASMA: Distributed Parallel Linear Algebra Software for Multicore Architectures, http://icl.cs.utk.edu/dplasma/.

[23] MAGMA (Matrix Algebra on GPU and Multicore Architectures), http://icl.cs.utk.edu/magma/.

[24] H. Kimura, M. Sato, Y. Hotta, T. Boku, D. Takahashi, Empirical study on reducing energy of parallel programs using slack reclamation by DVFS in a power-scalable high performance cluster, in: CLUSTER, 2006, pp. 1–10.

[25] M. Y. Lim, V. W. Freeh, D. K. Lowenthal, Adaptive, transparent frequency and voltage scaling of communication phases in MPI programs, in: SC, 2006, p. 107.

[26] P. Alonso, M. F. Dolz, R. Mayo, E. S. Quintana-Ortí, Improving power efficiency of dense linear algebra algorithms on multi-core processors via slack control, in: HPCS, 2011, pp. 463–470.

[27] P. Alonso, M. F. Dolz, F. D. Igual, R. Mayo, E. S. Quintana-Ortí, DVFS-control techniques for dense linear algebra operations on multi-core processors, CSRD 27 (4) (2012) 289–298.

[28] Intel Hyper-Threading Technology, http://www.intel.com/content/www/us/en/architecture-and-technology/hyper-threading/hyper-threading-technology.html.

[29] HPL - High-Performance Linpack Benchmark, http://www.netlib.org/benchmark/hpl/.

[30] T. Davies, C. Karlsson, H. Liu, C. Ding, Z. Chen, High performance linpack benchmark: A fault tolerant implementation without checkpointing, in: ICS, 2011, pp. 162–171.

[31] A. S. Wu, H. Yu, S. Jin, K.-C. Lin, G. Schiavone, An incremental genetic algorithm approach to multiprocessor scheduling, IEEE Transactions on Parallel and Distributed Systems 15 (9) (2004) 824–834.

[32] M. Weiser, B. Welch, A. Demers, S. Shenker, Scheduling for reduced CPU energy, in: OSDI, 1994, p. 2.

[33] A. Varma, B. Ganesh, M. Sen, S. R. Choudhury, L. Srinivasan, J. Bruce, A control-theoretic approach to dynamic voltage scheduling, in: CASES, 2003, pp. 255–266.

[34] K. Choi, R. Soma, M. Pedram, Fine-grained dynamic voltage and frequency scaling for precise energy and performance trade-off based on the ratio of off-chip access to on-chip computation times, in: DATE, 2004, pp. 18–28.

[35] Y. Taur, X. Liang, W. Wang, H. Lu, A continuous, analytic drain-current model for DG MOSFETs, IEEE Electron Device Letters 25 (2) (2004) 107–109.

[36] CPUSpeed, http://carlthompson.net/Software/CPUSpeed/.

[37] Watts up? Meters, https://www.wattsupmeters.com/.

[38] C. Liu, A. Sivasubramaniam, M. Kandemir, M. J. Irwin, Exploiting barriers to optimize power consumption of CMPs, in: IPDPS, 2005.

[39] C. Liu, J. Li, W. Huang, J. Rubio, E. Speight, X. Lin, Power-efficient time-sensitive mapping in heterogeneous systems, in: PACT, 2012, pp. 23–32.

[40] P. Alonso, M. F. Dolz, F. D. Igual, R. Mayo, E. S. Quintana-Ortí, Saving energy in the LU factorization with partial pivoting on multi-core processors, in: PDP, 2012, pp. 353–358.